DEFENCE
SOLUTIONS_

Statistics Text Classification

CASE STUDY

Methods
Analytics_

OUR CLIENT

Defence Digital Services (DDS) is a part of Strategic Command within the Ministry of Defence (MOD).

In the age of information warfare, Strategic Command plays a key role in the MOD, working across land, sea, air, space and cyber domains.

In 2020, Methods Analytics began working with DDS to create an Artificial Intelligence Centre of Expertise (AI CoE).

Our first project: develop an NLP solution to improve data validation for the compiling and reporting of the MOD health and safety annual statistics.

Strategic Command
Defence Digital

THE CHALLENGE WE FACED

Two analysts were responsible for collating and compiling health and safety data from across all MOD departments. This took four weeks a quarter.

The team had been asked to report health and safety outputs within MOD on a fortnightly basis. The team were struggling to meet this request using existing processes. Formed from seven distinct sources and with many fields purely free text, the data itself had varying levels of detail and completeness.

The manual interpretation of the data by the health and safety team also meant the data output was unreliable and unsuited to simply training a standard NLP model. The existing solution was also hard-coded to make classifications using MySQL.

From a more human perceptive, the NLP model and techniques would be new to the Health & Safety team. Comprehensive support would be required to ensure the team could confidently use the solution we recommended.

THE SOLUTIONS WE DEVELOPED

To prevent the value of the outputs degrading further, a machine learning solution using supervised learning was chosen. A more robust, adaptable and future-proof approach.

The model of success

Instead of manually interpreting the data, the team load in the collated unvalidated data to receive a classified output that reflects the existing process.

The integrated pipeline is modularised in line with best practice. These consist of pre-processing, utilities, feature engineering, machine learning and a main module to run and coordinate end-to-end.

This structure ensures it's not only simple for the health and safety analysts to use, but also that it can be easily understood when new functionality needs to be integrated.

Success of the productionised solution was defined as: 'outputs to be classified accurately 80% of the time'. But this does not take into account precision and recall. For this reason, we also included F-score as a metric for measuring the performance of the model.

The integration of the text classification model into existing processes has the power to remove most of the manual classification and all the SQL-based hardcoded classification.

The solution is a tiered system that first classifies any event in accordance with the existing nomenclature as 'sports and recreation', 'training', or 'normal duties'.

Secondary classification relates to the activity being undertaken (for example: if the first classification is sport, the secondary classification could be football).

There are also separate classifications running across all incidents which seek to identify the root cause of the event (for example: slip, trip, fall or struck by an object).

THE TEAM

By 2021, we had established the AI CoE, its operating model, capabilities and supporting functions.

The team consisted of over 25 specialists, engaging with stakeholders within DDS and delivering three AI projects alongside the development of standards, policy and guidance.

We worked closely with the Health & Safety team at DDS, considering their skills and needs in everything we did.

THE SUPPORT WE PROVIDED

One of the project's main aims was to empower the Health and Safety analysts – for them to use and maintain the model themselves.

A thorough training and knowledge sharing strategy was employed to bridge the relative inexperience of the analyst team with NLP techniques.

The programme included: regular sessions with the team to work through data quality issues, training sessions, and technical one-on-ones – from the fundamentals of coding in Python, to the specifics of maintaining and developing the model.

This ensured that the team were confident using the model and trusted the outputs.

Training content was documented in a comprehensive user guide, along with recommendations on how to improve the accuracy and efficiency of the initial solution.

THE GOALS WE ACHIEVED

1 An F-score of 1 represents the perfect balance between precision and recall, and for this project we aimed for >0.8.

2 For the finished product, the F-score was calculated at between 0.85 and 0.87 dependent of the health and safety first-tier category. This was achieved through refining the machine learning methods and regularly retraining the model.

3 To maintain and improve the F-score in the future, especially for the categories which did not have a wealth of data, the health and safety team were taught how to retrain the model based on updated validated data.

4 With this approach, the solution will continue to improve and the team can concentrate on developing the breadth and depth of the statistics and reporting.

We're Methods Analytics – a specialist data services company. We use data to help public and private sector clients solve complex problems and do good things.

Methods
Analytics_